WPE II Report - Prophet Inequalities with Limited Information - The Single Choice Problem and Beyond

Eshwar Ram Arunachaleswaran

May 2021

Abstract

The classical single choice prophet inequality problem has a gambler observing a sequence of rewards $X_1, X_2, .., X_n$, each of which is drawn from an associated probability distribution X_i D_i. At each step, the gambler must irrevocably decide whether to accept reward X_i (and stop) or reject it and move on to see reward X_{i+1}. The objective of the gambler is to obtain a reward that is, in expectation, within a constant (or reasonably small) factor of the reward obtained by an optimal offline player (i.e., a prophet). In the classical setting, the gambler has complete knowledge of the underlying distributions, and is known to have a 1/2-competitive algorithm. Similar results exist for more general settings, where the gambler has to select a set of rewards, subject to some conditions (for eg: selecting a set of rewards of size k). We study a new class of prophet inequalities, which do not assume that the gambler has full knowledge of the distributions. We look at the single sample setting, where the gambler only gets to look at one sample from each distribution before starting the sequence. In this setting, we see results from Azar et al. [AKW14] and Rubinstein et al. [RWW20], that prove competitive ratios in the single sample setting that are comparable to the corresponding best known competitive ratios for the classical version of the same problems. We also study the single choice problem when all rewards are independently drawn from a single, unknown distribution (Correa et al. [CDFS19]).

1 Introduction

In this report, we study a variant of the prophet inequality problem. In the classical single choice prophet inequality problem, n distributions $D_1, D_2 \cdots D_n$ are made known in advance to a gambler. Then, from $i = 1$ to n, a random variable V_i is drawn independently from D_i in each step and presented to the gambler. The gambler must decide irrevocably whether to accept V_i and stop the process or reject V_i and continue to the next step. The objective of the gambler is to come up with an algorithm that maximizes the expected reward (i.e., the value he accepts) obtained in this process, typically as a factor of the expected reward obtained by an offline player, known as the prophet, who sees all the draws before making her choice. This factor can be understood as an online competitive ratio for the gambler against an optimal offline adversary,i.e., the prophet.

The seminal result of Krengel,Sucheston and Garling [KS+77] showed the existence of a stopping rule for the single choice prophet inequality problem that guarantees the gambler an expected reward equal to at least half the expected reward of the prophet. This competitive ratio of 1/2 can be seen to be optimal through a simple example with two distributions [KS78].Sameul-Cahn [SC+84] showed how to obtain the same competitive ratio of 1/2 by using a threshold based stopping rule,i.e., the gambler stops and accepts any value greater than a threshold τ, that is fixed before seeing the draws from the distributions.

The secretary problem is a closely related, but incomparable, online optimization problem. We are given n non-negative numbers, in uniformly random order (every permutation of the numbers is equally likely), and once again, decisions must be made irrevocably whether to stop and accept any number or continue to the next number. The objective of the online player is to come up with a stopping rule to maximize the probability of selecting the maximum element in the sequence. The single choice secretary problem admits an elegant solution, in which the first $1/e$ fraction of the items are discarded and the first number that exceeds the maximum of the first $1/e$ fraction of the items is accepted. This solution guarantees a probability of $1/e$ of accepting the maximum element and is known to be the optimal solution to the secretary problem [F+89].

Recent work on prophet inequalities has focused on more general versions of the problem, in which the gambler's (and prophet's) objective is to select some maximum weight subset of the presented items, subject to some set system constraints, for example, selecting a subset of k out of the n items [Ala14], selecting an independent set of a matroid [KW12], selecting a set in a generic downward closed set system [Rub16], selecting a matching in a bipartite graph with online edge arrivals [GW19]. Similar generalizations exist for the secretary problem([KP09], [FSZ14], [DP08], [JSZ13]).

The variant of the prophet inequality problem we are interested in relaxes the assumption that the gambler has full knowledge of the distributions $\{D_i\}_{i\in[n]}$ from which the values are drawn. In this version of this problem, the only information that the gambler is assumed to have is access to a limited number of samples (all draws being independent) from each distribution. This problem was introduced by Azar et al. [AKW14], who gave the first results for a variety of problems (Single choice, k-choice Graphic Matroids, Laminar Matroids, Transerversal Matroids, Bipartite Matching) in this setting. A subset of these results are based upon a novel connection to a class of algorithms for the corresponding secretary problem 3, while others are based upon purpose built algorithms for those problems. Rubinstein et al. [RWW20] followed up by showing a $1/2$-competitive algorithm 2 for the single choice problem where the gambler has access only to a single sample from each distribution, improving upon the $1/4$-competitive algorithm from [AKW14] and closing the gap to the hardness bound of $1/2$. For the more general k-Choice Prophet Inequality problem, Azar et al. [AKW14] give a $1 - O\left(\frac{1}{\sqrt{k}}\right)$-competitive algorithm called the "Rehearsal Algorithm".

Another problem of interest for us is the single choice prophet inequality with identical distributions, known as the IID prophet inequality problem (IID standing, as is usual, for independent and identically distributed). The full information version of this problem has been well studied in many works, and is known to have an essentially optimal algorithm with competitive ratio 0.745 due to Correa et al. [CFH+17], matching an upper bound due to an impossibiliity result of Hill and Kertz [HK+82]. Correa et al. [CDFS19] considered a limited information of this problem where the gambler has no information at all about the distribution before the sequence begins. Observe that when all the values are drawn from the same distribution, any permutation of the observed values is equally likely, and hence the gambler can use the optimal algorithm for the secretary problem to obtain a competitive ratio of $1/e$. Surprisingly, Correa et al. [CDFS19] show that this algorithm is essentially optimal. Their result is based upon showing that for every algorithm, there is an infinite set $K \subset \mathbb{N}$ such that when $V_1, V_2 \cdots V_i$ are supported on K, the decision to stop at V_i when $V_i > \max_{j\in[i-1]} V_j$ is independent of $V_1, V_2 \cdots V_i$ (Lemma 1). This property, referred to as value-obliviousness, is established by employing the infinite version of Ramsay's Theorem [Ram30], and can then be used to connect the guarantees of algorithms for this problem to guarantees for algorithms for the secretary problem (in a reduction of sorts). Additionally, the same work also studies an intermediate version of the problem where the gambler gets to see a few samples from this unknown distribution. While the same hardness result is shown to hold for $o(n)$ samples, they show an improved ratio of $1 - 1/e \approx 0.632$ with $n-1$ samples and that $O(n^2)$ suffice to get arbitrarily close to the upper bound of 0.745. Rubinstein et al. [RWW20] further strengthened this result by showing that the gambler can get arbitrarily close to 0.745-competitive with $O(n)$ samples.

We present and analyze a cross section of results from the three papers [AKW14], [RWW20] and [CDFS19].

1.1 Further Related Work

There is a rich vein of connections between prophet inequalities and mechanism design established by the work of Hajiaghayi et al. [HKS07] and Chawla et al. [CHMS10]. This connection is inspired by the problem of picking a set of buyers in online fashion based on some prior information about their valuations. Many of the results referred to in this report translate into posted price mechanisms in the corresponding mechanism design problem environments.

The work of Azar et al. [AKW14] was in part inspired by results in the same flavor in mechanism design([DRY15] and [HR09]), that explored revenue maximization based on limited information about buyer's priors.

1.2 Summary of Results

We tabulate the results for a collection of prophet problems under different set system constraints (which we refer to as the environment). In Table 1, we compare the best known algorithms, under full information as well as limited information, as well as hardness results (i.e., upper bounds on the competitive ratios). Note that upper bounds for the full information setting are also upper bounds for the limited information setting. The number of samples refers to the number of sets of samples, where each set consists of one sample from each distribution.

Environment	Full Information Lower Bounds	Limited Information Lower Bounds	Upper Bounds
Single Choice	1/2-competitive [KS+77]	1/2-competitive [RWW20]	1/2 (Full Information) [KS78]
k-Choice	$\left(1-\frac{1}{\sqrt{k+3}}\right)$-competitive [Ala14]	$1-O\left(\frac{1}{\sqrt{k}}\right)$-competitive [AKW14]	$\left(1+\frac{1}{\sqrt{512k}}\right)^{-1}$ (full information) [HKS07]
Bipartite Matching on (V,E)	1/3-competitive [GW19]	$\frac{1}{6.75}$ ($\min\{\lvert E\rvert, d^2\}$-samples where d is max degree) [AKW14]	1/2.25 (full information) [GW19]
Matroid Independent Set	1/2-competitive [KW12]	$O(\log\log($ rank $))$-competitive (combining the result of [FSZ14] and [AKW14])	1/2 (full information) [KW12]
IID Prophet	0.745-competitive [CFH+17]	$1/e$ (no samples) ; $1-1/e$ (n-1 samples) [CDFS19] ; $0.745-\varepsilon$ ($O_\varepsilon(n)$ samples) [RWW20]	0.745 (full information) [HK+82] ; $1/e+\delta$ ($o(n)$ samples) [CDFS19]

Table 1: Comparison of Results Across Different Environments

1.3 Outline of Report

Single Choice: In Section 3, we contrast one of the simplest algorithms for the full information single choice prophet inequality problem (from Kleinberg et al. [KW12]) with the algorithm of Rubinstein et al. [RWW20] for the single sample version of the same problem.

Connections between Secretary and Prophet Problem: In Section 4, we study the result of Azar et al. [AKW14] that show how to utilize a certain class of algorithms, called order-oblivious algorithms, for the secretary problem to construct algorithms with the same competitive ratio for the equivalent version (same environment) of the single sample prophet inequality problem.

Bipartite Matchings: In Section 5, we study the algorithm of Azar et al. [AKW14] that ensures a competitive ratio of 1/6.75 for the problem of finding a bipartite matching with $O(|E|)$ samples. This algorithm illuminates how to exploit the availability of multiple samples to establish useful independence properties between certain events. Of interest is also the setting of asymmetric thresholds for different edges in a combinatorial environment, with similarities to techniques used in mechanism design for multi-dimensional sequential posted price environments [CHMS10].

IID Prophet Inequalities: In Section 6, we study the results of Correa et al. [CDFS19] about the limited information version of the IID Prophet Inequality problem. The main result we study is the hardness result presented in Section 6.2, which shows that the gambler cannot do better than a competitive ratio of $1/e + \delta$ for the unknown IID prophet inequality problem. We also briefly touch upon the algorithmic results that show how to beat this lower bound and close the gap to the upper bound of 0.745 with access to $\theta(n)$ samples.

2 Notation and Preliminaries

2.1 Online Environment

Given an environment $\mathcal{I} = \{U, \mathcal{J}\}$ where U is a universe set of indices $U = [n]$ and $\mathcal{J}$ is a set of subsets of U, the online selection problem takes as input an ordered sequence of positive real numbers, referred to as values, $V = (v_1, v_2, \cdots v_n)$ with each v_i being revealed in a single step. As each element v_i is revealed, the algorithm $\mathcal{A}$ must irrevocably decide whether or not to include v_i in a set A of accepted elements, subject to the constraint that $A \subset \mathcal{J}$ throughout the process. Without loss of generality, we can assume that $\mathcal{J}$ is a downwards closed set system. Throughout the report, we use capitalization (V_i) to refer to a random variable and lower case (v_i) to refer to the realization of the random variable (V_i).

2.2 Prophet Inequalities

The prophet inequality problem is a special case of the online selection problem, where each value V_i is drawn independently from some distribution D_i. All these distributions are assumed to exclusively be supported on non-negative numbers. Let $\mathcal{D}$ be the product distribution $D_1 \times D_2 \cdots D_n$. We can alternatively consider $v \sim \mathcal{D}$ as the input. In the classical version of the problem, the offline player, referred to as the gambler, is completely aware of each distribution D_i, but has no control over the order in which the values are presented. The objective is for the algorithm to pick a set whose value is competitive, in expectation, with respect to the optimal set picked by an offline adversary, the eponymous prophet. In particular, let $\mathrm{OPT}(V)$ represent $\max_{J \in \mathcal{J}} \sum_{i \in J} V_i$. Then, we say that algorithm $\mathcal{A}$ selecting set $A(V)$ has a prophet inequality with competitive ratio α if:

$$\mathbb{E}_{v\sim\mathcal{D}}[\sum_{i \in A(V)} V_i] \geq \alpha . \mathbb{E}_{v \sim D}[\mathrm{OPT}(V)]$$

We refer to $\sum_{i \in A(v)} v_i$ as the reward obtained by the algorithm.

Implicitly, we also allow the expectation on the left hand side to be taken over random choices made by the algorithm $\mathcal{A}$ as well as the randomness in the draws from the distribution.

Let $V_{\max}$ represent the random variable equal to the max of $\{V_i \sim D_i\}$. The expected reward of the prophet is $\mathbb{E}[V_{\max}]$ for the single choice prophet inequality.

Limited Information Prophet Inequality Problems: While the classical prophet inequalities problem assumes that the gambler has full knowledge of the distributions, the limited information version of these problems gives the algorithm l independent draws from each distribution as the only information about them before the sequence begins. When $l = 1$, the problem is called the single sample prophet inequality problem. These samples are indexed by S_j^i with realizations s_j^i where the upper index is over the sample set, and the lower index over the element in the universe set that it corresponds to. The upper index is dropped when it is clear from context.

Note. *For our convenience, we can assume without loss of generality that all samples and values are distinct. To justify this, we associate an independent draw from Uniform*$(0, 1)$ *for each draw and use this value (distinct with probability one) to break ties.*

2.3 The Secretary Problem

The secretary problem is also a special case of the online selection problem, where the values $v_1, v_2, \cdots v_n$ are shown in uniformly random order. Formally, n underlying values $\{v_i\}_{i \in [n]}$ are permuted uniformly at random and shown in a sequence $v_{i_1}, v_{i_2} \cdots v_{i_n}$. The objective of the secretary problem is to come up with an algorithm that maximizes the probability of picking the optimal set in $\mathcal{J}$. The worst case (over all sequences of numbers) probability associated with an algorithm is referred to as its competitive ratio α.

$$\Pr[\mathcal{A}_s \text{ selects } \mathrm{OPT}(v)] \geq \alpha$$

Again, we allow this probability to be over both the randomness in the order of the elements as well as random choices of the algorithm.

Observe that maximizing the probability of choosing a maximum value set (uniqueness is WLOG) also implies a competitive bound when taking the ratio of the expected value of the set chosen by the algorithm and the expected value of the maximum value feasible set.

In the most basic version of the secretary problem, the single choice problem, we wish to maximize the probability of picking the maximum element in the sequence.

3 A $1/2$-Competitive Algorithm for the Single Sample Single Choice Problem

The single sample single choice prophet inequality problem is the problem of selecting a single element (where $\mathcal{J}$ consists of all singletons sets and the empty set) where the gambler is only given access to a single sample drawn independently (of the value draws) from each distribution. The result of Rubinstein et al. [RWW20] show that a simple threshold based algorithm gives a prophet inequality with competitive ratio $1/2$.

In this section, we look at upper bounds for the competitive ratio of any algorithm for this problem, and then contrast a simple pre-existing algorithms for the full information version of the problem and the algorithm of [RWW20]. Note that the Single Sample problem is weakly harder than the full information version and any algorithm for this setting automatically translates into an algorithm for the full information version.

3.1 Upper Bound for the Single Choice Prophet Inequality

First, we look at a simple instance to see that $1/2$ is the best possible ratio that can be guaranteed. Let D_1 be a distribution that puts its entire probability mass on 1 and let D_2 put probability mass ε on $\frac{1}{\varepsilon}$ and the rest on 0 where $\varepsilon > 0$ is a very small positive real number, i.e., $\varepsilon << 1$. Note that $\mathbb{E}[v_1] = \mathbb{E}[v_2] = 1$. We assume that the algorithm has full knowledge of the distributions, making this an upper bound instance for the traditional single choice prophet inequality problem as well. Note that the algorithm gets no new information on observing $v_1 \sim D_1$. Therefore, we can restrict our analysis to two classes of algorithms - those that always accept the first value, and those that never accept the first value (and presumably accept the second value, since it is always no worse than not accepting the second value). Observe that both algorithms have an expected reward of 1. In contrast, we show that the prophet has an expected reward of $2 - \varepsilon$. We can look at the following two events that partition the probability space, the first event is $v_2 = \frac{1}{\varepsilon}$ and the second event is $v_2 = 0$. The first event occurs with probability ε and gives reward $\frac{1}{\varepsilon}$ to the prophet while the second event occurs with probability $1 - \varepsilon$ and gives reward 1 to the prophet, thus giving us the desired bound.

3.2 A Simple Algorithm for the Full Information Version [KW12]

Algorithm $\mathcal{A}_{KW}$ [KW12]: Set threshold $\tau = \mathbb{E}[V_{\max}]$ and accept the first value V_i greater than the threshold.

Theorem 1. *$\mathcal{A}_{KW}$ is a 1/2-competitive algorithm for the single choice prophet inequality problem.*

To see the proof of the algorithm. We split the reward of the algorithm R_A into two parts $R_A = R_T + R_E$. For the first part R_T, note that any value picked by the algorithm is at least as large as the threshold, we set $R_T = \min\{R_A, \tau\}$. The second part is the "excess", defined as $R_E = \max\{R_A - \tau, 0\}$. Note that $\mathbb{E}[R_A] = \mathbb{E}[R_T] + \mathbb{E}[R_E]$. Let F be the event that the algorithm accepts any of the observed values in the sequence and let $\overline{F}$ be its complementary event. Clearly, $\mathbb{E}[R_T] = \tau.\Pr[F]$. Since this is a single choice problem, the algorithm stops when it accepts a value. Conditional upon the algorithm reaching a value V_i, this value contributes $\mathbb{E}[\max\{V_i - \tau, 0\}]$ towards the quantity $\mathbb{E}[R_E]$ (we are implicitly breaking up $\mathbb{E}[R_E]$ into the expected contribution from each value in the sequence using the linearity of expectation). Thus, we have:

$$\mathbb{E}[R_E] = \sum_{i\in[n]} \mathbb{E}[\max\{V_i - \tau, 0\}]\Pr[V_i \text{ is reached by } \mathcal{A}_{KW}]$$

Let F_i be the event that value V_i is reached by the algorithm $\mathcal{A}_{KW}$. Note that all events $\{F_i\}_{i\in[n]}$ happen if event $\overline{F}$ occurs. Thus, we get the bound:

$$\begin{aligned}
\mathbb{E}[R_E] &\geq \sum_{i\in[n]} \mathbb{E}[\max\{V_i - \tau, 0\}].\Pr[\overline{F}] \\
&\geq \mathbb{E}[\max_{i\in[n]} V_i - \tau].\Pr[\overline{F}] \\
&= \tau.\Pr[\overline{F}]
\end{aligned}$$

Thus, the expected reward of the algorithm is at least $\tau(\Pr[\overline{F}] + \Pr[F]) = \tau = \mathbb{E}[X_{\max}]/2$. This gives the desired competitive ratio of 1/2.

3.3 Algorithm of Rubinstein et al. [RWW20]

Let any algorithm observe samples $S_i \sim D_i$ for all i before starting the sequence for the online selection problem.

Algorithm $\mathcal{A}_{\mathbf{sssc}}$ [RWW20] : Set threshold $\tau = \max_i S_i$ and only accepts the first value V_i that is greater than the threshold.

Theorem 2 ([RWW20])**.** *The algorithm $\mathcal{A}_{sssc}$ has a competitive ratio of* 1/2*.*

The key observation made in [RWW20] is that the samples and values can equivalently be viewed in the following manner. First, two draws Y_i, Z_i are made from the distribution D_i (note: while it can be assumed that the draws are independent, the analysis works even if they are arbitrarily correlated), WLOG, let $Y_i > Z_i$. Then we independently toss an unbiased coin B_i, if heads (i.e., $b_i = 1$), $S_i = Y_i$ and $V_i = Z_i$ and if tails, we flip the assignment. We will see that for every draw of $\{y_i, z_i\}_{i\in[n]}$, in expectation over the coin flips $\{B_i\}_{i\in[n]}$, the given algorithm gives a 1/2-competitive prophet inequality.

Consider the sorted order on the set $\{x : x = y_i \text{ or } x = z_i\}$. This order looks like $y_{i_1} > y_{i_2} \cdots > z_{j_1} \cdots > z_{j_n}$. Relabel this sequence as $X = X_1 > X_2 \cdots > X_{2n}$. Each X_i corresponds to some distribution index j. Let i^* be the smallest index in $[2n]$ such that there exists a index $j^* < i^*$ such that both X_{i^*} and X_{j^*} correspond to the same distribution index. Therefore, $X' = X_1, X_2, \cdots X_{i^*-1}$ are all of the form Y_{i_1} to $Y_{i_{i^*-1}}$. Importantly, this means that the coin flips assigning the elements of X' to samples or values are all independent. The intuition behind the competitive ratio is as follows. Observe that the prophet picks the first value in the sequence X and the algorithm picks a value (if at all one exists) that is at least larger than the first sample in the sequence. To make our analysis easier, let us restrict the reward obtained by the algorithm by assuming that it only gets the smallest value (if one exists) that is larger than the threshold (i.e., the largest sample).

Note that X_1 is a value with probability $1/2$ (based on the associated coin flip) and is always picked by the prophet given that it is a value. To guarantee that X_1 is picked by the algorithm, X_1 must be a value and X_2 a sample - this happens with probability $1/4$ (using the independence of the first few coin flips). Similarly, X_2 is picked by the prophet with probability $1/4$ (X_1 is a sample and X_2 is a value) and picked by the algorithm with probability $1/8$ (X_1, X_2 are values and X_3 is a sample). This reasoning can be extended all the way up to the index $i^* - 2$ since all coin flips $\{b_{i_l}\}_{l\in[i^*-1]}$ are independent. This happens with probability $\frac{1}{2^{i^*-1}}$. X_{i^*-1} is chosen by the algorithm only when the first sample in the sequence is X_{i^*}, which in turn only happens when the first $i^* - 1$ coin flips result in the corresponding X_i being set as a value. For X_{i^*}, note that the prophet picks this quantity only when the first $i^* - 1$ coin flips result in the corresponding X_i being set as a sample. Observe that the first sample as well as the first value in the sequence both appear within the index i^*, thus we do not have to consider the possibility of either the prophet or algorithm selecting X_j with $j > i^*$. Putting this together, we get the following bounds on the rewards of the prophet R_P and the algorithm R_A:

$$
\begin{aligned}
\mathbb{E}[R_P] &= \left(\sum_{i=1}^{i^*-1} \frac{X_i}{2^i}\right) + \frac{X_{i^*}}{2^{i^*-1}} \\
\mathbb{E}[R_A] &\geq \left(\sum_{i=1}^{i^*-2} \frac{X_i}{2^{i+1}}\right) + \frac{X_{i^*-1}}{2^{i^*-1}} \\
&= \left(\sum_{i=1}^{i^*-1} \frac{X_i}{2^{i+1}}\right) + \frac{X_{i^*-1}}{2^{i^*-2}}
\end{aligned}
$$

Comparing the two expressions leads to the 1/2-competitive ratio.

3.4 Comparison

Note that the algorithm $\mathcal{A}_{KW}$, for the full information setting, required some non trivial information about the statistics of the distributions. In particular, it required us to know the expected value of the maximum of the set composed of a single draw from each distribution. As noted by [Luc17], this algorithm is robust to small errors in the estimation, the expected reward drops by ε if there is a 2ε- error in the estimate for the expected value of the maximum. However, estimating this quantity still requires some knowledge of the distributions or a good number of samples depending upon the properties of the distributions. In contrast, the single sample algorithm $\mathcal{A}_S$, achieves the same bound using just one sample from each distribution. However, it is interesting to note that the analysis of [RWW20] does not go through when we only have *a single sample of the random variable* V_{max} even though we only use the maximum value of the samples in our algorithm. The analysis of $\mathcal{A}_S$ uses these samples to explicitly reason about all possible selections by the algorithm (and the prophet). In contrast, the analysis of $\mathcal{A}_{KW}$, the full information setting algorithm, does not explicitly analyze which value is selected by the algorithm.

4 Reducing Prophet Inequalities to the Secretary Problem

In this section, we present the results of [AKW14] that show how to use a certain class of algorithms for the secretary problem to generate the same competitive ratio for the corresponding prophet inequality problem.

We say that algorithm $\mathcal{S}$ is an order-oblivious algorithm for the secretary problem if it satisfies the following conditions.

1. $\mathcal{S}$ must behave in the following manner- $\mathcal{S}$ picks a threshold index k before starting the sequence (potentially using random bits) and only observes the first k values $\{v_{i_1}, v_{i_2} \cdots v_{i_k}\}$ in the sequence (and does not accept any of these values).

2. $\mathcal{S}$ assumes only that the set A is a uniformly random subset of size k of the set $\{v_i\}_{i\in[n]}$ of n values, while proving the competitive ratio. Note that this is a strict subset of the conditions that are given to the algorithm.

To illustrate these conditions, let us consider a simple order oblivious algorithm with a competitive ratio of $1/4$ for the single choice secretary problem. The algorithm observes the first k values without accepting where $k = \text{Bin}(n, 1/2)$, sets a threshold $\tau = \max_{l\in[k]} v_{i_l}$ and accepts any value in the second half that is greater than the threshold. To lower bound the probability that the algorithm stops at the maximum value, consider the two largest elements v_i and v_j, with $v_i < v_j$. We consider the first k elements to be a uniformly random subset of size k. This set of k elements can be equivalently considered as tossing k independent unbiased coins to decide whether each element v_l goes into the first k elements. Thus, with probability $1/4$, v_i is in the first k elements and v_j is not, implying that the maximum element is chosen with probability at least $1/4$.

Theorem 3 ([AKW14])**.** *Any α-competitive order-oblivious algorithm $\mathcal{A}_S$ for the secretary problem in environment $\mathcal{I} = \{U, \mathcal{J}\}$ yields a α-competitive algorithm $\mathcal{A}_P$ for the corresponding single sample prophet inequality problem in the same environment.*

The intuition behind this result is that we can use the samples to recreate the conditions required for the order-oblivious algorithm $\mathcal{A}_S$ with the key observation being that samples and values are statistically equivalent.

We now describe the algorithm $\mathcal{A}_P$. We construct a sequence $X = X_1, X_2 \cdots X_n$ to run algorithm $\mathcal{A}_s$ on, and show that X simultaneously satisfies the preconditions the algorithm $\mathcal{A}_S$ requires and is similar enough to our real sequence V for the guarantees of $\mathcal{A}_S$ to translate into a prophet inequality.

Algorithm $\mathcal{A}_P$ based on Order-Oblivious Algorithm for the Secretary Problem in Environment $\mathcal{I} = \{U, \mathcal{J}\}$

1. (Offline) Let k be the threshold chosen by algorithm $\mathcal{A}_S$.

2. (Offline) Let π be a uniformly random permutation on $[n]$. Pass $S_{\pi(1)}, S_{\pi(2)}, \cdots S_{\pi(k)}$ as the first k elements of sequence X where $S = (S_1, S_2, \cdots S_n)$ is the sample set.

3. (Online) Run the algorithm $\mathcal{A}_S$ on sequence X, which is constructed as the values V_i are revealed.

4. (Online) As each value V_i is revealed, add V_i as the next element in sequence X if $i \notin \{\pi(1), \pi(2), \cdots \pi(k)\}$, otherwise ignore V_i.

First, we observe that the accepted set A consists purely of elements in $\{X_{k+1}, X_{k+2} \cdots X_n\}$. Additionally, since both problems share the same feasible set system, the accepted set A is present in the superset $\mathcal{J}$ of feasible sets.

Observe that the expected reward of the prophet is $\mathbb{E}[\mathrm{OPT}(V)]$. By construction of the sequence X, the random variables $\mathrm{OPT}(X)$ and $\mathrm{OPT}(V)$ are statistically equivalent, implying that their expectations are equal. However, since X satisfies the conditions required by the analysis of $\mathcal{A}_S$, we know that $\mathbb{E}[\sum_{i\in A(X)} X_i] \geq \alpha.\mathbb{E}[\mathrm{OPT}(X)]$. Observe that the left hand side of the expression is exactly the expected reward obtained by the algorithm, leading to the desired competitive ratio for $\mathcal{A}_P$.

The above theorem results in the following corollary.

Corollary 1.

1. *For general matroid constraints, there exists a $O(\log\log(rank))$-competitive algorithm for the single sample prophet inequality problem, based upon the secretary algorithm of Feldman et al [FSZ14].*

2. *For graphic matroids, there exists a $1/8$-competitive algorithm for the single sample prophet inequality problem, based upon the secretary algorithm of Korula and Pal [KP09].*

3. *For transversal matroids, there exists a $1/16$-competitive algorithm for the single sample prophet inequality problem, based upon the secretary algorithm of Dimitrov and Plaxton [DP08].*

4. *For laminar matroids, there exists a $1/16$-competitive algorithm for the single sample prophet inequality problem, based upon the secretary algorithm of Jaillet, Soto and Zenklusen [JSZ13].*

5 $1/6.75$-Competitive Algorithm for Bipartite Matching

Consider the problem environment, where we have an underlying bipartite graph $G = (A \cup B, E)$. The set $\mathcal{J}$ consists of matchings of G. In each step, we get the weight V_e on edge e, drawn from the corresponding distribution D_e. First, we show the result using d^2 samples for maximum degree d graphs, and then mention how to achieve the same competitive ratio for general bipartite graphs with $|E|$ samples. We get d^2 sets of samples $S^1, S^2 \cdots S^{d^2}$, where each sample set $S^i = \{S^i_1, S^i_2, \cdots S^i_{|E|}\}$ is a set of independent draws from the distribution of each edge.

Arbitrarily number the edges incident on each vertex. Consider any edge $e = (u, v)$, where e is the i-th edge incident on u and the j-th edge incident on v. We define the index of this edge, represented by z_e to be $d.i + j$. This indexing has two important properties, first that all indices are in $\{1, 2, \cdots d^2\}$, second that no two edges sharing an endpoint have the same index.

For each edge, we define an associated threshold τ_e. Consider any vector of edge weights $w = (w_1, w_2 \cdots w_e)$. We know there exist efficient deterministic algorithms for maximum weight bipartite matchings. Let $x_e(w)$ be 1 if such an algorithm include edge e in the bipartite matching on the graph G with edge weights e. We define a threshold $T_e(w_{-e})$ where w_{-e} refers to the set of weights of all edges excepting e.

$$T_e(w_{-e}) := \inf\{w_e : x_e(w_e, w_{-e}) = 1\}$$

Intuitively, this threshold is the price of admission for edge e, if all the other edge weights are known. Since we do not know the other edge weights in advance, we use a sample to set the actual threshold as $\tau_e := T_e(s^{z_e}_{-e})$. The intuition for this threshold is straightforward - the expected contribution of edge e to the offline optimum is $\mathbb{E}[V_e \geq T_e(V_{-e})].\Pr[x_e(V_e, V_{-e}) = 1]$. Note that $\Pr[x_e(V_e, V_{-e}) = 1] = \Pr[x_e(V_e, S^{z_e}_{-e}) = 1]$ since S^{z_e} and V are statistically identical. Using the same reasoning, $\mathbb{E}[V_e|V_e \geq \tau_e] = \mathbb{E}[V_e|V_e \geq T_e(v_{-e})]$.

We show how to use these threshold along with simple randomized selection to achieve the target competitive ratio.

Algorithm $\mathcal{A}_B$ for selecting a bipartite matching:

1. Initialize a set of accepted elements A to ϕ.
2. In the e-th step, we see the value of edge e, which is $V_e = v_e$.
   a Flip a coin B_e independently such that $B_e = 1$ w.p. 1/3 and $B_e = 0$ w.p 2/3.
   b Discard edge e if $B_e = 0$ and move to the next edge.
   c If $B_e = 1$, add edge e to A if $A \cup \{e\}$ is a matching and $V_e \geq \tau_e$.

Theorem 4. *The algorithm $\mathcal{A}_B$ is 1/6.75-competitive for the bipartite matching prophet inequality problem on graphs with maximum degree d.*

We have already seen that the contribution of any edge e to the offline optimum OPT is $\mathbb{E}[V_e|V_e \geq \tau_e].\Pr[V_e \geq \tau_e]$. Thus, using the linearity of expectation, it suffices to prove that the contribution of edge e to the expected reward of the algorithm, denoted by W_e is at least $(1/6.75).\mathbb{E}[V_e|V_e \geq \tau_e].\Pr[V_e \geq \tau_e]$.

Let C_e denote the event that edge e is accepted by algorithm $\mathcal{A}_B$. Let A_e denote the state of the set A when value on edge e is revealed to the algorithm. We have $W_e = \mathbb{E}[V_e|C_e].\Pr[C_e]$. Let X_e be the event that $C_e = 1$, Y_e be the event that $V_e \geq \tau_e$ and Z_e be the event that $A_e \cup \{e\}$ is a matching. Clearly, $C_e = X_e \cap Y_e \cap Z_e$. Additionally, X_e happens with probability 1/3 and independent of the other two events as well as the draw V_e,. Thus, we have:

$$W_e = (1/3)\Pr[Y_e \text{ and } Z_e].\mathbb{E}[V_e|Y_e \text{ and } Z_e]$$

However, Y_e and Z_e are not independent - this is because Y_e depends upon the sample S^{z_e}, which in turn might also affect other edges whose values are revealed before V_e, in turn affecting the probability that $A \cup \{e\}$ is a matching.

The right lens to look at the event $Y_e \cap Z_e$ is through the following two events E_1 and E_2. Let edge e have endpoints l and r. E_1 is the event that no other edge incident on l is selected by $\mathcal{A}_B$ and E_2 is the event that no other edge incident on r is selected by $\mathcal{A}_B$. They key property that we exploit is that E_1 and E_2 are each independent of Y_e as well as the draw V_e, since two edges that share an edge have thresholds set based on different samples. Observe that the occurrence of E_1 and E_2 guarantee that event Z_e will happen. Thus, we get:

$$\begin{aligned} W_e &\geq (1/3)\Pr[Y_e \text{ and } E_1 \text{ and } E_2].\mathbb{E}[V_e|Y_e \text{ and } E_1 \text{ and } E_2] \\ &= (1/3).Pr[Y_e].\Pr[E_1 \text{ and } E_2].\mathbb{E}[V_e|Y_e] \\ &= (1/3)Pr[V_e \geq \tau_e].\Pr[E_1 \text{ and } E_2].\mathbb{E}[V_e|V_e \geq \tau_e] \end{aligned}$$

Thus, it suffices to show a bound of 4/9 on the quantity $\Pr[E_1 \text{ and } E_2]$ to obtain the desired result.

Consider the probability space of any sample S^i- we see that the total probability of selecting an edge e as part of a maximum weight matching on the graph weighted with S^i is $\Pr[S^i_e \geq T_e(S^i_e, S^i_{-e})]$. Additionally, we know that the probability of selecting an edge incident on a particular vertex v is at most 1, by the definition of a matching. Using the fact that all samples and the values are identical random variables, we get : (Recall that $e = (l, r)$)

$$\sum_{e:e=(l',r)} \Pr[Y_{e'}] \leq 1$$

$$\sum_{e:e=(l,r')} \Pr[Y_{e'}] \leq 1$$

Additionally, for any edge e', since $C_{e'} = \cap Y_{e'} \cap Z_{e'}$, we know that $\Pr[C_{e'}] \leq \Pr[X_{e'} \cap Y_{e'}] = 1/3 \Pr[Y_{e'}]$. Thus, we have:

$$\sum_{e:e=(l',r)} \Pr[C_{e'}] \leq \frac{1}{3}$$

$$\sum_{e:e=(l,r')} \Pr[C_{e'}] \leq \frac{1}{3}$$

Thus, we have $\Pr[E_1], \Pr[E_2] \leq \frac{1}{3}$. It is not hard to show that $\Pr[E_1|E_2] \geq \Pr[E_1]$ and thus we get $\Pr[E_1 \cap E_2] \geq 4/9$, completing the proof. Note that we chose a probability of $1/3$ to set $C_e = 1$ since $x = 1/3$ maximizes $x(1-x)^2$ in $[0,1]$.

To modify the algorithm to use $|E|$ samples for general bipartite graphs - we change the indexing and just set index of edge e to be e, we still retain the property that two edges incident on the same vertex have different indices.

6 IID Prophet Inequalities

Another variation of the prophet inequality problem is the unknown IID Prophet Inequality problem, introduced and studied by [CDFS19]. Recall that in this problem, all distributions D_i are identically the same distribution D but the gambler has no knowledge of this distribution before the sequence begins.

6.1 A $1/e$-comeptitive Algorithm for the Unknown IID Prophet Inequality Problem

We see a simple reduction to the single choice secretary problem to obtain a $1/e$-competitive algorithm for the unknown IID Prophet Inequality Problem.

Theorem 5. *There exists a $1/e$-competitive algorithm for the unknown IID Prophet Inequality problem.*

Proof. Consider any draw of values $V = \{v_1, v_2, v_3, \cdots v_n\}$ and any permutation π on $[n]$. Observe that any permutation of the values $V = \{v_{\pi(1)}, v_{\pi(2}, v_{\pi(3)} \cdots v_{\pi(n)}\}$ is equally likely to be observed. Thus, conditioned on any set of n values being drawn, every permutation of this set is equally likely to be observed. Thus, running the optimal $1/e$-competitive algorithm [F+89] for the single choice secretary problem on the observed sequence guarantees that the maximum value is selected with probability at least $1/e$. This gives the desired bound on the expected reward of the algorithm. ☐

6.2 A $1/e + \delta$ Upper Bound for the Unknown IID Prophet Inequality Problem [CDFS19]

More interestingly, Correa et al. [CDFS19] showed that the single choice secretary based approach outlines above is essentially optimal, by showing that any algorithm for this problem (in a canonical form that preserves the competitive ratio) has a "difficult" instance on which it cannot do much better than $1/e$.

Before stating this result, we introduce some special notation .

We say that an algorithm for this problem, referred to a as a stopping rule $\mathbf{r}$ (since in the single choice problem, we either stop and accept a revealed value or we continue), has stopping time τ if it stops at index τ. To model the fact that the underlying distribution D is not known to the gambler, we describe $\mathbf{r}$ as a collection of functions $\mathbf{r}_i : \mathbb{R}^i_+ \to [0,1]$ where $p = \mathbf{r}_i(v_1, v_2, \cdots v_i)$ is the probability with which the stopping rule accepts v_i at the i-th step, conditioned upon seeing values $V_1 = v_1$, $V_2 = v_2 \cdots$, $V_i = v_i$ and not having stopped at any of the steps from 1 to $i-1$. Note that implicitly, the functions $\mathbf{r}_i$ also depend on n. Wherever we wish to make this clear, we explicitly refer to stopping rules as a function of n, i.e., as $\mathbf{r}(n)$.

Theorem 6 ([CDFS19]). *For any $\delta > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$ and any stopping rule $\mathbf{r}(n)$ with stopping time τ, there exists a distribution F, unknown to the stopping rule, such that $\mathbf{r}(n)$ has competitive ratio upper bounded by $(1/e + \delta)$ for the (unknown) IID prophet inequality problem with distribution F,i.e.,*

$$\mathbb{E}[V_\tau] \leq (1/e + \delta).V_{max}$$

Before outlining the proof of this theorem, we observe that the same hardness result can also be shown to hold for the relaxed version of the problem where the algorithm gets to see $o(n)$ samples before the sequence begins. To provide some intuition, assume that some stopping rule $\mathbf{r}$ can do better than $1/e+\delta$ in this relaxed setting. Now, for the original unkown IID prophet inequality problem, consider the following algorithm - there is no stopping at the first n' values, these are considered to be "samples' and we run the algorithm $\mathbf{r}$ for the relaxed problem on the remaining $n-n'$- long sequence. Let us set $n' = o(n)$, thus $n-n' = (1-o(1)n$ and the expected maximum value in this shortened sequence is at least $(1-o(1)n$ the expected maximum of the original sequence. Thus, we show the existence of a stopping rule with competitive ratio $(1/e\delta).(1-o(1)$ for the unknown IID prophet inequality problem, contradicting the above theorem.

The proof of Theorem 6 is based on two ideas. The first idea is that $1/e$ is the best possible competitive ratio for the single choice secretary problem. The second idea is that, for any stopping rule (modified to a canonical form), there is a specific distribution for which this rule is "value-oblivious",i.e., it cares only about the relative comparisons between the observed values and not the values themselves. Additionally, if the largest realized value dominates all the other values, then this stopping rule essentially behaves like an algorithm for the single choice secretary problem. The following lemma formalizes this intuition.

Definition 1. *Let $\varepsilon > 0$, $V \subset \mathbb{N}$. We say that a stopping rule $\mathbf{r}(n)$ is value-oblivious on V if for all $i \in [n]$, there exists $\{q_i\}_{i\in[n]}$ such that for all distinct $v_1, v_2 \cdots v_i \in V$ with $v_i > \max_{j\in[i-1]} v_j$, we have the property that $\mathbf{r}_i(v_1, v_2 \cdots v_i) \in [q_i - \varepsilon, q_i + \varepsilon]$.*

Lemma 1. *Consider any $\varepsilon > 0$, given a stopping rule $\mathbf{r}'(n)$ with competitive ratio α, then there exists a stopping rule $\mathbf{r}(n)$ with competitive ratio α and an infinite set $V \subset N$ such that $\mathbf{r}(n)$ is value-oblivious on the set V.*

Given this lemma, we see a proof of Theorem 6.

Let $\varepsilon = 1/n^2$. Consider any stopping rule $\mathbf{r}(b)$ with competitive ratio α. By Lemma 1, we know there exists a stopping rule $\mathbf{r}'(b)$ with competitive ratio α and an infinite set $V \subset \mathbb{N}$ such that $\mathbf{r}(n)$ is value-oblivious on V.

We build the distribution F in the following manner. Since V is an infinite set, we can assume there are natural numbers $s_1, s_2, \cdots s_{n^3}, m$ such that $m \geq n^3 \max_i s_i$. We set probability masses on these elements as follows - for all $i \in [n^3]$. The probability mass set by F on eachc V_i is defined as follows:

$$V_i = \begin{cases} s_1, & \text{with probability } 1/n^3(1-1/n^2) \\ s_2, & \text{with probability } 1/n^3(1-1/n^2) \\ & \vdots \\ s_{n^3}, & \text{with probability } 1/n^3(1-1/n^2) \\ m, & \text{with probability } 1/n^2 \end{cases}$$

We observe some immediate consequences of this construction.

1. With high probability, $V_1, V_2 \cdots V_n$ are all distinct.
2. The expectation of $V_{\max}$ is dominated by the expected contribution of the maximum element, i.e., $\mathbb{E}[V_{\max}] \geq \frac{m}{n}$.
3. The expectation of V_τ (the reward of the stopping rule) is dominated by the contribution of the event that m appears among $V_1, V_2 \cdots V_n$ and that $V_1, V_2 \cdots V_n$ are distinct, and additionally the contribution

of this event is itself dominated by the possibility of the algorithm stopping at m,i.e., if τ is the stopping time, then

$$\mathbb{E}[V_\tau] \leq (1/n). \Pr[V_\tau = m | V_{\max} = m \text{ and } V_1, V_2 \cdots V_n \text{ are distinct}].m + o(1/n).m$$

Let $p_1 := \Pr[V_\tau = m | V_{\max} = m$ and $V_1, V_2 \cdots V_n$ are distinct]

We only need an upper bound of $1/e + o(1)$ on the probability p_1 to obtain the desired bound on the competitive ratio α. Recall that our stopping rule $\mathbf{r}'(n)$ is ε-value oblivious (since all values observed are distinct and restricted to the set V). Additionally, since we are only interested in the probability of getting the maximum value, we can assume that $\mathbf{r}'(n)$ never accepts a value unless it is the largest value seen so far, without any loss of generality.

Consider any stopping rule $\mathbf{r}''(n)$ that is 0 value-oblivious on the set V, with stopping time τ''. If $\Pr[V_{\tau''} = m | V_{\max} = m$ and $V_1, V_2 \cdots V_n$ are distinct$] > 1/e$, we can use $\mathbf{r}''(n)$ as an algorithm for the secretary problem, contradicting $1/e$ being the optimal competitive ratio for that problem.

Now, we see some intuition as to why we can obtain a similar bound on the probability p_1 even though $\mathbf{r}'(n)$ is only ε value-oblivious. For all $i \in [n]$, we know that for distinct $v_1, v_2 \cdots v_i \in V$ with $v_i > \max_{j \in [i-1]} v_j$, we have the property that $\mathbf{r}'_i(v_1, v_2 \cdots v_i) \in [q_i - \varepsilon, q_i + \varepsilon]$. Construct a 0 value-oblivious stopping rule $\mathbf{r}''(n)$ with stopping time τ'' by setting $\mathbf{r}''_i(v_1, v_2 \cdots v_i) = q_i$. Through a coupling argument and the use of the union bound, it can be seen that

$$\Pr[V_\tau = m | V_{\max} = m \text{ and } V_1, V_2 \cdots V_n \text{ are distinct}] \leq \Pr[V_{\tau''} = m | V_{\max} = m \text{ and } V_1, V_2 \cdots V_n \text{ are distinct}] + n.\varepsilon$$

Since $n.\varepsilon = 1/n = o(1)$, we get the desired result.

6.2.1 Proof of Lemma 1

The heart of the hardness result for the unknown IID prophet inequality problem lies in Lemma 1, which shows that any stopping rule in a canonical form cannot make use of the actual observed values when they are drawn from a suitable infinite set. This is the essential ingredient that allows us to force the stopping rule to use any more information than an algorithm for the single choice secretary problem.

First, we set up a canonical form for stopping rules and argue why converting any stopping rule into this canonical form preserves the competitive ratio.

Definition 2. *A stopping rule $\mathbf{r}(n)$ is said to be order-oblivious if for all $i \in [n]$, any distinct sequence $v_1, v_2, ...v_i$ and any permutation π of $[i-1]$, $\mathbf{r}_i(v_1, v_2, \cdots v_i) = \mathbf{r}_i(v_{\pi(1)}, v_{\pi(2)}, \cdots v_{\pi(n-1)}$.*

Lemma 2. *If there exists a stopping rule $\mathbf{r}(n)$ with competitive ratio α, then there exists an order-oblivious stopping rule $\mathbf{r}'(n)$ with the same competitive ratio and such that for all $i \in [n]$ $\mathbf{r}'(n)$ never selects V_i if $\mathbf{r}(n)$ never selects V_i.*

Proof Sketch of Lemma 2 First, we describe how to construct $\mathbf{r}'(n)$ given a stopping rule $\mathbf{r}(n)$. For every $i \in [n]$, let $\sim_i$ be the equivalence relation on $\mathbb{R}^i_+$ such that $(v_1, v_2 \cdots v_n) \sim_i (w_1, w_2 \cdots w_n)$ if $w_1, w_2 \cdots w_{i-1}$ is a permutation of $v_1, v_2 \cdots v_{i-1}$ and $v_i = w_i$. We say that a stopping rule $\mathbf{r}$ arrives at index i through $(v_1, v_2 \cdots v_n)$ if $V_1 = s_1, V_2 = s_2 \cdots V_i = s_i$, $(s_1, s_2 \cdots s_i) \sim_i (v_1, v_2, \cdots v_n)$ and $\mathbf{r}$ does not stop at any index $j \in [i-1]$.

$$\mathbf{r}'_i(v_1, v_2, \cdots v_n) := \Pr[\mathbf{r} \text{ stops at index } i \; | \mathbf{r} \text{ arrives at index } i \text{ through } (v_1, v_2, \cdots v_n)_{\sim_i} \;]$$

By definition, $\mathbf{r}'(n)$ is order-oblivious. Additionally, $\mathbf{r}'(n)$ never selects V_i if $\mathbf{r}(n)$ never selects V_i. We will give some intuition as to why $\mathbf{r}'(n)$ has the same competitive ratio as $\mathbf{r}(n)$. We can think of $\mathbf{r}'(n)$ as uniformly distributing the probability of $\mathbf{r}(n)$ accepting v_i, having arrived at index i through $(v_1, v_2, \cdots v_n)_{\sim_i}$ over all the different sequences in this equivalence class. We can show by induction over i that for all $i \in [n]$, we have $\Pr[\mathbf{r}$ arrives at index i through $(v_1, v_2, \cdots v_n)_{\sim_i}\ ] = \Pr[\mathbf{r}'$ arrives at index i through $(v_1, v_2, \cdots v_n)_{\sim_i}]$. Let $\mathbf{r}(n)$ and $\mathbf{r}'(n)$ have stopping times τ and τ' respectively. Consider the contribution of the event that v_i is accepted after arriving at index i through $(v_1, v_2, \cdots v_n)_{\sim_i}$. It is not hard to show that the expected contribution of this event is the same for both stopping rules.

Thus, starting with $\mathbf{r}(n)$, we now have a order-oblivious stopping rule $\mathbf{r}'(n)$.

Now, the final step, which is proving the value-obliviousness of $\mathbf{r}'(n)$, can be done index by index - a stopping rule is said to be (ε, i) value-oblivious if it is value-oblivious on some infinite set K_i upto the i-th step. We prove that $\mathbf{r}'(n)$ is value-oblivious by induction. Additionally, the corresponding infinite sets $K_1, K_2 \cdots K_n$ will have the property that $K_1 \supseteq K_2 \cdots \supseteq K_n$.

For the base case, we consider $i = 0$, and we know that every stopping rule trivially always continues to the next value, with $V_0 = \mathbb{N}$.

To do the induction step, we call upon the infinite hypergraph version of Ramsay's Theorem [Ram30].

Theorem 7 ([Ram30])**.** *Let H be a d-uniform infinite complete hypergraph whose edges coloured with c colours. Then, H must have a monochromatic d-uniform infinite complete sub-hypergraph.*

Consider the induction hypothesis for index i. We have assumed there exists infinite set $K_{i-1} \subseteq N$ such that $\mathbf{r}'(n)$ is $(\varepsilon, i-1)$ value-oblivious on K_{i-1}. Consider the i-uniform infinite complete hypergraph H with vertex set K_{i-1}. We partition the unit interval $[0, 1]$ into $\lceil \frac{1}{2\varepsilon} \rceil$ intervals - $[0, 2\varepsilon), [2\varepsilon, 4\varepsilon) \cdots [1 - 2\varepsilon, 1)$. We associate a colour C_i with each of these intervals, in total we have $c = \lceil \frac{1}{2\varepsilon} \rceil$ colours. Consider any hyperedge $(v_1, v_2, \cdots v_i)$ where $v_i \geq \max_{j \in [i-1]} v_j$. We know that there exists a unique natural number u such that $\mathbf{r}'_i(v_1, v_2 \cdots v_i) \in [(2u-1)\varepsilon - \varepsilon, (2u-1)\varepsilon + \varepsilon)$ - we colour this hyperedge with the colour corresponding to this interval. Using Ramsay's Theorem(Theorem 7), we know there exists a monochromatic d-uniform infinite complete sub-hypergraph of H. Call the vertices of this hypergraph K_i. Observe that $K_i \subseteq K_{i-1}$, K_i is an infinite set, and since all hyperedges have the same colour - corresponding to some interval $[(2u-1)\varepsilon - \varepsilon, (2u-1)\varepsilon + \varepsilon)$, we conclude that $\mathbf{r}'(n)$ is (ε, i) value-oblivious on K_i, successfully extending the induction hypothesis.

6.3 Beating $1/e$ with $O(n)$ Samples

Although we cannot expect to do better than a competitive ratio of $1/e$ with $o(n-1)$ samples, we already have some reason to believe that this state of affairs changes when the gambler has access to $O(n)$ samples. In particular, with exactly n samples, the gambler can run the algorithm of Rubinstein et al. [RWW20] to achieve a competitive ratio of $1/2$.

First, we see a simple algorithm that achieves a $1 - 1/e$-competitive ratio when the gambler is given access to $n(n-1)$ samples. Then, we provide some intuition as to how the same result can be achieved with $n - 1$ samples.

Consider the algorithm that takes samples $\{S_i = s_i\}_{i=1}^{n(n-1)}$ as input.

$1 - 1/e$ Algorithm $\mathcal{A}_{n(n-1)}$ for IID Prophet Inequality with $n(n-1)$ Samples

1. Let the samples be $\{S_i = s_i\}_{i=1}^{n(n-1)}$. For each $i \in [n]$, allocate a distinct set of $n-1$ samples, and let τ_i be the maximum value of this set.
2. Accept the i-th value if $V_i > \tau_i$, let τ be the stopping time if V_i is accepted.

Conditioned upon the algorithm reaching index i, we make two key observations -

1. The probability of accepting V_i is $\frac{1}{n}$, since V_i has to be the maximum of n independent draws from the same distribution. Additionally, this event is independent of the probability of accepting value V_j conditioned upon reaching V_j for $j \neq i$- this is because we use a disjoint set of samples to set

the threshold for each index. Thus, the unconditional probability of reaching V_j is $\left(1-\frac{1}{n}\right)^{j-1}$ for all $j \in [n-1]$.

2. The expected value of V_i conditioned upon the algorithm accepting V_i is $\mathbb{E}[V_i|V_i \geq \tau_i] = \mathbb{E}[V_{\max}]$, since V_i is accepted only if it is the maximum of n independent draws from the same distribution.

Thus, we can write the expected reward of the algorithm as:

$$\begin{aligned}
\mathbb{E}[V_\tau] &= \sum_{i=1}^{n} \Pr[\ \mathcal{A}_{n(n-1)} \text{ reaches value } V_i\ ].\Pr[\ \mathcal{A}_{n(n-1)} \text{ accepts value } V_i \mid \mathcal{A}_{n(n-1)} \text{ reaches value } V_i\ ].\mathbb{E}[V_i|V_i \geq \tau_i] \\
&= \sum_{i=1}^{n} \left(1-\frac{1}{n}\right)^{i-1}.\frac{1}{n}.\mathbb{E}[V_{\max}] \\
&= \left(1-\left(1-\frac{1}{n}\right)^{n}\right)\mathbb{E}[V_{\max}] \\
&\geq \left(1-\frac{1}{e}\right)\mathbb{E}[V_{\max}]
\end{aligned}$$

Now consider the IID Prophet problem where the gambler has $n-1$ samples. The previous algorithm required $n-1$ fresh samples at the time where it arrived at the i-th value. The modification that Correa et al. [CDFS19] is to pick a random subset of size $n-1$ from $\{S_1, S_2, \cdots S_{n-1}, V_1, V_2, \cdots V_{i-1}\}$ and use the maximum value in this set as a threshold on arriving at value V_i.

The following lemma shows that this modification preserves the conditions that we use to analyze the algorithm $\mathcal{A}_{n(n-1)}$.

Lemma 3. *Let $U_i := \{S_1, S_2, \cdots S_{n-1}, V_1, V_2, \cdots V_{i-1}\}$. Conditioned upon the modified algorithm arriving at value V_i, the distribution of the set U_i is identical to $n+i-2$ independent draws from the underlying distribution D.*

This lemma is not at all obvious, given that the algorithm arriving at this step (rejecting the first $i-1$ values) seems to implicitly induce some correlations between the samples and values in the set U_i. We note that the proof is via induction on i.

With this lemma, we can still preserve the key independence properties that allowed us to argue the $1-1/e$-competitive ratio guarantee for algorithm $\mathcal{A}_{n(n-1)}$.

Theorem 8 ([CDFS19])**.** *There exists a $1-1/e$-competitive algorithm for the IID Prophet Inequality problem with the gambler given access to $n-1$ samples.*

We do not provide a proof of the above lemma or the resulting theorem in this report.

We note that this result may potentially lead to an interesting corollary about the unknown IID Prophet Inequality problem when the underlying distribution D has some particular properties.

Consider the following algorithm for the unknown IID Prophet problem - ignore the first half of the sequence, and use the first $n-1$ values as samples for the second half of the sequence. Let K_l represent the random variable that represents the maximum of l independent draws from D. We know that for all positive support distributions, $\frac{\mathbb{E}[K_{n/2}]}{\mathbb{E}[K_n]} \geq 1/2$. This trivially gives our algorithm a lower bound of $(1-1/e)/2 \approx 0.31$. However, $\frac{\mathbb{E}[K_{n/2}]}{\mathbb{E}[K_n]} \geq 1/2$ is only tight for certain distributions, for eg., the class of distributions used in the hardness result where the expectation of a single draw depended heavily upon the appearance of a single dominant value. For distributions where this ratio approaches 1 as $n \to \infty$, the unknown IID prophet problem can be solved with a competitive ratio of $(1-1/e)-\varepsilon$ for large n, significantly improving upon $1/e$. This property is for instance true for the uniform distribution on $[0,1]$, where $\mathbb{E}[K_l] = \frac{l}{l+1}$.

Finally, we briefly discuss closing the gap to $0.745v - \varepsilon$ competitive ratio using $O_\varepsilon(n)$ samples. The algorithm of Rubinstein et al. [RWW20] for this problem builds upon the algorithm of Correa et al. [CFH+17] for the full information version of the IID Prophet problem. They first show that the statistical information about D used by the algorithm of Correa et al. [CFH+17] can be approximately constructed using $O_\varepsilon(n)$ samples. Then, they show that this algorithm in [CFH+17] is sufficiently robust to still achieve the same competitive ratio (upto a small additive loss) using the approximated statistics about the underlying distribution D.

7 Future Directions and Open Problems

A significant open problem is that of finding a constant factor competitive algorithm for the single sample matroid prophet problem. For the k-choice prophet problem, it would be interesting to nail down an exact, rather than an asymptotic competitive factor, particularly for small k. For the IID prophet problem, we would like to know the limits of what competitive factor can be achieved with exactly $n-1$ samples. Another closely related problem is to identify classes of distributions for which the hardness result (Theorem 6) can be bypassed to achieve better than $1/e$ competitive factors.

References

[AKW14] Pablo D Azar, Robert Kleinberg, and S Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1358–1377. SIAM, 2014.

[Ala14] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[CDFS19] José Correa, Paul Dütting, Felix Fischer, and Kevin Schewior. Prophet inequalities for iid random variables from an unknown distribution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 3–17, 2019.

[CFH+17] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 169–186, New York, NY, USA, 2017. Association for Computing Machinery.

[CHMS10] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.

[DP08] Nedialko B Dimitrov and C Greg Plaxton. Competitive weighted matching in transversal matroids. In *International Colloquium on Automata, Languages, and Programming*, pages 397–408. Springer, 2008.

[DRY15] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[F+89] Thomas S Ferguson et al. Who solved the secretary problem? *Statistical science*, 4(3):282–289, 1989.

[FSZ14] Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple o (log log (rank))-competitive algorithm for the matroid secretary problem. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1189–1201. SIAM, 2014.

[GW19] Nikolai Gravin and Hongao Wang. Prophet inequality for bipartite matching: merits of being simple and non adaptive. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 93–109, 2019.

[HK+82] Theodore P Hill, Robert P Kertz, et al. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability*, 10(2):336–345, 1982.

[HKS07] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.

[HR09] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009.

[JSZ13] Patrick Jaillet, José A Soto, and Rico Zenklusen. Advances on matroid secretary problems: Free order model and laminar case. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 254–265. Springer, 2013.

[KP09] Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *International Colloquium on Automata, Languages, and Programming*, pages 508–520. Springer, 2009.

[KS+77] Ulrich Krengel, Louis Sucheston, et al. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4):745–747, 1977.

[KS78] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

[KW12] Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.

[Luc17] Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.

[Ram30] Frank P Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 2(1):264–286, 1930.

[Rub16] Aviad Rubinstein. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 324–332, 2016.

[RWW20] Aviad Rubinstein, Jack Z Wang, and S Matthew Weinberg. Optimal single-choice prophet inequalities from samples. *Innovations in Theoretical Computer Science*, 2020.

[SC+84] Ester Samuel-Cahn et al. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, 12(4):1213–1216, 1984.